Clock Signal 01: AI Co-Scientists and the Research Clock

What changes when the definition of "hard" keeps moving?

Prepared by Patrick Halford – Gaoithe Advisory (www.gaoitheadvisory.com)

Field	Value
Scan Run Date	**14 June 2026**
Baseline Window	15 May 2026-14 June 2026
Scan Type	FP10 Technology Clock Watch - Clock Signal 01
Evidence Boundary	Public evidence only; European FP10 assumptions tested against relevant global capability movement.

Executive Summary

The research clock moved again in the last 30 days.

The live signal is not that AI replaces scientists. That is the cheap version. The real signal is that AI is starting to change the operating tempo around scientists, R&D teams, corporate labs, SMEs, RTOs, spinout offices and consortia.

Google DeepMind published and launched Co-Scientist as a multi-agent research partner built with Gemini, designed to generate, debate, refine and rank scientific hypotheses. Google also launched Gemini for Science tools, including Hypothesis Generation, Computational Discovery, Literature Insights and Science Skills, with explicit links to Co-Scientist, AlphaEvolve, NotebookLM and specialist life-science databases. These tools are being positioned not as search assistants, but as workbench infrastructure for parts of the scientific method.

That matters because FP10 is being designed for 2028-2034, while the research process itself is already moving. The Commission has proposed a EUR 175bn Horizon Europe programme for 2028-2034, and Europe is also building RAISE as a virtual institute to pool compute, data, talent and funding for AI in science.

The problem is timing. FP10 calls, consortia and work packages are slow machines. AI co-scientists, scientific agents and global R&D loops are fast machines. When those machines meet, the old assumptions start to rot.

> *Is this still hard for the same reason?*

What needed a large consortium six months ago may now need a sharper technical operator. What looked too early for spinout may now be commercially testable. What looked like a strong work package may now be stale. What looked like an SME weakness may become SME leverage if cheaper tools reduce capability barriers.

The first conclusion is blunt: FP10 will not only be shaped by European policy. It will be pulled by global AI model progress, scientific-agent systems, open-source tooling, corporate R&D adoption, product launches, compute access, regulation and capital movement.

The second conclusion is practical: this cannot be handled as a one-off read. It needs a rescan rhythm. The baseline keeps moving. Universities, RTOs, corporates, SMEs, spinout teams and consortia need to see which assumptions changed before budgets, partners and work packages get locked around yesterday's version of difficulty.

1. The Signal

AI co-scientists are moving from research curiosity to early research infrastructure.

Google's Co-Scientist is built as a coalition of specialised agents that generate hypotheses, cluster research directions, critique ideas, rank them through an idea tournament, refine the strongest directions and produce a research proposal for human review. Google says the system integrates web search and specialist databases such as ChEMBL and UniProt, and can use specialised models such as AlphaFold in selected collaborations.

That is a shift from AI as helper to AI as structured scientific reasoning layer.

The most important change is not raw intelligence. It is workflow compression:

Research activity	Old assumption	New pressure
Literature synthesis	Manual and slow	AI can structure large bodies of literature faster
Hypothesis generation	Human-led and narrow	AI agents can generate and debate many candidate hypotheses
Proposal shaping	Periodic and labour-heavy	AI can assist earlier framing and comparison
Computational testing	Limited by human bandwidth	Agentic systems can generate and test more variants
Scientific review	Human-only bottleneck	AI-assisted critique and peer-review tooling is emerging

This does not remove scientific judgement. It raises the value of scientific judgement. Bad teams will use these tools to generate more noise. Strong teams will use them to ask better questions earlier.

2. Latest 30-Day Signal Check

Baseline window: 15 May-14 June 2026

Signal	Freshness	Evidence category	Why it matters
Google DeepMind published Co-Scientist and opened interest for Hypothesis Generation.	Last 30 days	Observed evidence	Moves AI co-scientist from concept into accessible experimental tooling.
Google launched Gemini for Science tools: Hypothesis Generation, Computational Discovery, Literature Insights and Science Skills.	Last 30 days	Observed evidence	Shows the toolset widening from one co-scientist system to a research workbench.
Google says enterprise-grade scientific and industrial R&D tools are already in private preview with partners including BASF, Daiichi Sankyo, Bayer Crop Science and U.S. National Labs.	Last 30 days	Observed evidence from company source	Important for corporate R&D clock: this is not only academic.
AutoScientists paper submitted 27 May 2026. It describes decentralised agent teams for long-running computational scientific experimentation, with reported gains across biomedical ML, GPT training optimisation and protein fitness prediction.	Last 30 days	Preprint evidence	Shows movement toward long-running, multi-agent research workflows.
EinsteinArena paper submitted 9 June 2026. It reports an agent-native platform where agents work on open problems with verifiers, leaderboards and discussion; the paper reports 12 new state-of-the-art results as of May 2026.	Last 30 days	Preprint evidence	Signals collective agent research, not just isolated single-agent work.
A paper submitted in late May argues AI research agents may narrow scientific exploration by staying close to seed literature.	Last 30 days	Preprint evidence	This is the cold shower. AI agents may improve local search while weakening broader scientific exploration.
Anthropic announced Claude Fable 5 and Mythos 5 on 9 June and then posted a 12 June update suspending access. Its own launch page claims strong scientific-research capability and longer autonomous work, but access volatility is the bigger signal.	Last 30 days	Observed company evidence	Shows both capability pressure and governance/access risk.
The European Commission recorded June 2026 AI governance moves, including launch of the AI Act Practical Forum and Scientific Panel on 1 June and a Tech Sovereignty Package on 3 June.	Last 30 days	Observed policy evidence	Europe is tightening the governance and sovereignty clock while global capability accelerates.

Reading: The last 30 days were not quiet. They showed movement on three fronts at once: capability, workflow tooling and governance. That combination matters more than any single announcement.

3. Why This Matters for FP10

FP10 is being designed as a 2028-2034 programme. The Commission's stated ambition is a larger, faster and more impactful Horizon Europe, with a proposed EUR 175bn budget.

But the research process is changing before the programme starts.

This creates a design problem:

FP10 design assumption	Pressure from AI co-scientists
Calls can be designed around stable technical baselines	Technical baselines may shift before calls close
Work packages can remain valid for multi-year projects	Work package assumptions may decay inside the project
Consortia are assembled around known capability gaps	AI may change what capability is scarce
Commercialisation comes later	Market testing may need to move earlier
Evaluation can focus on excellence, impact and implementation	Adaptive capacity may become a fourth practical test

FP10 should not only fund research topics. It should fund research systems that can update their assumptions while the science is moving.

4. The Different Clocks

FP10 consortia do not run on one clock. They are a stack of misaligned clocks.

Actor	Clock problem
Universities	Research cycles, PhDs, publications and grant incentives move slower than AI-enabled discovery tools.
RTOs	Validation and testbeds may become more important, but only if they can respond faster.
Corporates	Product roadmaps may shift faster than consortium milestones.
SMEs	AI can give SMEs leverage, but cash and customer pressure still force brutal prioritisation.
Spinouts	AI can pull spinout timing forward, but global comparators can also make a spinout obsolete faster.
Consortia	A consortium can be formally on track while its original logic is already old.
FP10 system	Calls, missions and evaluation cycles lag global capability movement.

The strategic risk sits in the mismatch.

5. Global Dependency Loops

The scan is FP10-facing. The pressure is global.

Loop	Current signal	What this means
Research loop	Co-Scientist, AutoScientists and EinsteinArena show movement in hypothesis generation, agent teams and verified open problems.	Research groups need to compare their current methods against new agentic workflows.
Product loop	Gemini for Science moves the signal toward usable tools and science workbenches.	Tools are becoming part of the baseline, not special projects.
Capital loop	Private-preview use by corporates and labs signals commercial seriousness, even before broad adoption.	Corporates will not wait for FP10 cycles if internal R&D tools move faster.
Supply chain loop	Compute, data access and specialist databases matter more as research systems become agentic.	Capability depends on data, compute and tool access, not only expertise.
Regulation loop	EU AI governance moved again in June 2026.	Compliance and access will shape who can use frontier tools inside consortia.
Customer loop	Corporate R&D users are being pulled toward faster testing and product relevance.	Commercialisation timing moves earlier.
Talent loop	Researchers who can supervise AI-agent workflows become more valuable.	Training and operating model gaps become real bottlenecks.
Geopolitical loop	Europe is pairing AI-in-science ambition with sovereignty and governance moves.	FP10 must handle global capability without becoming dependent or slow.

6. Self-Improvement Pressure Layer

Runaway self-improving AI should not be treated as the base case.

The practical signal is narrower and more useful:

Level	Status	Read
1. Assisted improvement	Already here	AI helps with writing, coding, literature work, analysis and comparison.
2. Process improvement	Now visible	Gemini for Science and Co-Scientist show AI being embedded into research workflow steps.
3. Agentic improvement	High watch priority	AutoScientists and EinsteinArena show agents organising, testing, learning from failures and reusing other agents' ideas.
4. Recursive improvement	Strategic scenario, not base case	There are signs of AI improving parts of AI/research workflows, but not enough to treat full recursive self-improvement as the core assumption.

The useful framing:

> *The issue is not whether AI becomes fully self-improving in one dramatic leap. The issue is whether enough parts of research and R&D become AI-improved that the overall clock accelerates.*

7. Assumption Decay

A project can stay on budget and still drift away from reality.

Assumption area	Old baseline	New signal	What may have changed	Practical implication
Budget	More money needed for more discovery	AI can reduce cost of search, synthesis and computational testing	Some pilots may need less money, but more compute/data discipline	Recheck budget logic before bid or rescope
Resources	Scarce human analyst/researcher time	Agents can handle parts of the workload	Human bottleneck shifts to supervision and judgement	Train teams to manage AI-assisted workflows

Assumption area	Old baseline	New signal	What may have changed	Practical implication
Capability	Large institutions hold most capability	SMEs and spinouts can rent or access advanced tools	Smaller teams may punch above their old weight	Reassess partner roles
Partnership need	Broad consortia needed to cover gaps	AI may reduce some gaps and expose new ones	Some partners become redundant; others become critical	Run partner freshness review
Time	Research path assumed multi-year	Hypothesis and computational testing may compress	Commercial testing may move earlier	Bring market review forward
Commercialisation	Later-stage activity	AI makes early option testing cheaper	Spinout and licensing paths may surface earlier	Add commercialisation rescan every quarter
Spinout readiness	Technology must mature first	AI can compress validation, investor preparation and market mapping	Some spinouts become testable sooner	Run spinout timing review
Risk	Technical risk dominates	Market, competitor and access risk rise	A good project can be overtaken globally	Add global competitor watch

8. The Hardness Map

Area	What used to be hard	What may now be easier	What may now be harder	What this changes
Literature review	Reading, structuring, comparing papers	Large-scale synthesis and gap spotting	Knowing whether synthesis is shallow or biased	Need human review discipline
Hypothesis generation	Creating enough plausible paths	Generating and ranking many hypotheses	Avoiding crowded, derivative ideas	Need novelty checking
Experiment design	Planning and prioritising tests	Computational experiment generation	Real-world lab validation still hard	RTOs become more valuable
Simulation/modelling	Coding and variant testing	Agentic code generation and parallel testing	Compute and data quality constraints	Budget shifts from labour to infrastructure
Coding/tool-building	Specialist developer time	AI-assisted coding and debugging	Production-grade reliability	SMEs gain leverage but still need engineering judgement
Validation	Slow, expensive, physical	Earlier computational narrowing	Wet lab, field and industrial validation remain bottlenecks	RTO/testbed role strengthens
Data access	Finding and cleaning data	Better connectors and specialist databases	Data rights, quality and interoperability	Data strategy becomes central
Commercial testing	Expensive market discovery	Faster market mapping and prototype narratives	Buyers still need trust and proof	Earlier customer conversations
Spinout formation	Slow packaging of tech and market story	Faster evidence packs and investor materials	Global comparators move faster too	Spinout windows need active watching
Consortium design	Partner search and work package logic	Faster capability mapping	Stale partner assumptions	Need clock stress test before submission

9. Actor-by-Actor Impact

Universities

Universities gain from faster hypothesis generation, literature synthesis and early technical exploration. But they also face a harder benchmark. If global labs and private players use agentic research systems faster, university teams can be outpaced while still publishing good work.

The spinout implication is sharp: TTOs and science ventures need to rescan technologies more often. Some assets become commercial earlier. Others become less defensible.

RTOs

RTOs may become more important, not less. AI can generate and narrow hypotheses. It does not automatically validate them in industrial environments. The RTO role shifts toward fast validation, testbeds, standards, safety and evidence.

The stronger RTO role is no longer only to help translate research. It is to stop stale assumptions entering delivery.

Corporates

Corporate R&D teams will respond faster than universities when the link to roadmap, product risk or market pressure is clear. Google's enterprise-preview positioning around scientific and industrial R&D tools is a clear signal that the corporate clock is moving.

For corporates in FP10 consortia, the question becomes: does the work package still match the product clock?

SMEs

SMEs gain leverage when tools lower the cost of analysis, coding, experimentation and market testing. But they are also exposed. Larger companies can use the same tools with more data, customers and capital.

An SME's consortium role should be rescanned for two things: new leverage and new squeeze.

University Spinouts

Spinouts need a timing review. AI may help produce stronger evidence packs, faster market maps, better technical comparisons and sharper investor materials. But the global competitor clock also speeds up.

A spinout that looked early may now be testable. A spinout that looked defensible may now be crowded.

Consortia

Consortia are at high risk of hidden assumption decay. The visible project plan may still look clean. The underlying logic may not.

The most exposed areas are work packages, partner roles, validation timing, market assumptions and IP defensibility.

FP10 System

FP10 should consider whether adaptive capacity becomes part of project quality. Static excellence is not enough if the baseline moves during the grant.

The practical question for FP10 is: Can a project update its assumptions without breaking its governance?

10. Rescan Requirement

This scan must be repeated monthly or quarterly.

Monthly is better for this signal because the tooling and model layer is moving fast. Quarterly may be enough for slower institutional buyers.

Each rescan should ask:

Rescan question	Purpose
What changed since the last scan?	Keep baseline live
Which assumptions are now stale?	Prevent wasted money
What became easier or cheaper?	Update delivery logic
What became harder or scarcer?	Expose risk
Which partners became more important?	Fix consortium design
Which partners became less essential?	Reduce dependency
Which work packages need review?	Protect project value
Which spinout paths moved forward?	Capture earlier value
Which global competitor/product changed the baseline?	Avoid blind spots

The value does not come from a one-off scan. It comes from repeated baseline checks.

11. Decision Triggers

Trigger	Who should care	Why it matters	Possible response
Work package assumption decay	Consortium coordinators	Project can be on track but stale	Review assumptions before the next decision point
Partner clock mismatch	Universities, RTOs, corporates	Partners move at different speeds	Recheck roles against current capability
Spinout timing shift	TTOs, science ventures, VCs	Commercial timing may move earlier or later	Reassess commercial testing timing
Corporate R&D roadmap pressure	Corporates	Product clock may outrun project clock	Compare work packages with roadmap pressure
SME capability uplift	SMEs, coordinators	SMEs may now carry more weight	Reassess SME role, leverage and exposure
RTO validation bottleneck	RTOs, funders	Validation becomes the choke point	Prioritise validation and testbed readiness
Global product or competitor shift	All actors	External movement can kill assumptions	Monitor external product, competitor and research movement
AI-enabled research acceleration	Universities, RTOs	Research tempo changes	Review workflows against AI-enabled methods
Missing adaptive capacity	FP10 bid teams	Static plans weaken	Test whether assumptions can be updated within governance

12. Practical Response Areas

Response area	Who should review it	What should be checked	Why it matters
Research-clock briefing	University and RTO leadership	Whether workflows reflect AI-enabled hypothesis generation, synthesis and testing	Leaders need a clear view before FP10 assumptions harden
Consortium clock review	Coordinators and bid teams	Whether work packages, milestones and roles still fit the current baseline	A consortium can look complete while its logic is already stale
Horizon portfolio rescan	Universities, innovation offices and science ventures	Which live projects have moved assumptions around cost, timing, partners or commercialisation	Existing projects may contain earlier commercial or spinout options
Spinout timing review	TTOs, science ventures, VCs and founders	Whether AI-enabled workflows change readiness, defensibility, validation timing or investor comparators	Some spinouts become testable earlier; others become crowded faster
Corporate R&D baseline review	Corporate R&D and innovation teams	Whether research still matches roadmap pressure and global competitor movement	Corporate product clocks can outrun consortium clocks

13. Watchlist for the Next Scan

Watch item	Why it matters	Source types to monitor
New AI co-scientist systems	Shows whether Google is alone or a market forms	AI labs, arXiv, Nature, Science
Validated scientific use cases	Separates evidence from demos	Peer-reviewed papers, lab case studies
University adoption	Shows when this becomes normal research infrastructure	University news, lab pages, grants
Corporate R&D adoption	Shows commercial seriousness	Corporate R&D blogs, partnerships
SME tool leverage	Shows whether smaller firms gain capability	SME product releases, open-source tools
Spinout acceleration	Shows commercial timing changes	TTO news, VC deals, startup launches
FP10 / EU policy movement	Connects global signal to European programme design	Commission, Parliament, Council, NCPs
Global competitor movement	Tests European assumptions	US, UK, China, Japan, Korea, India, Australia sources
Product launches	Changes buyer and market baseline	Company releases, developer platforms
Capital flows	Shows where investors believe timing has moved	VC funding, M&A, corporate venture

14. Closing Interpretation

The first issue of the FP10 Technology Clock Watch shows a clear signal: the research clock is no longer only human, institutional and grant-driven. It is being pulled by AI-enabled hypothesis generation, agentic experimentation, scientific workbenches, global corporate R&D and policy pressure.

The mistake would be to turn this into another AI trend report.

The strategic value is in the rescan.

The question to put in front of universities, RTOs, corporates, SMEs, spinout teams and consortia is simple:

> *Which of your assumptions were true when the project was written, but are no longer true now?*

In FP10, the danger is not only choosing the wrong research topic.

The danger is continuing to fund, staff and partner around yesterday's idea of what is hard.

Selected Source Base

The report is based on public evidence. Source links are included below for review and updating in the next rescan.

- Google DeepMind - Co-Scientist: a multi-agent AI partner to accelerate research
- Google DeepMind - Gemini for Science: AI experiments and tools for a new era of discovery
- European Commission - Horizon Europe 2028-2034: twice bigger, simpler, faster and more impactful
- European Commission - AI in Science / RAISE
- European Commission - European approach to artificial intelligence
- arXiv - AutoScientists
- arXiv - EinsteinArena
- arXiv - AI research agents and exploration breadth
- Anthropic - Claude Fable 5 and Mythos 5